PCI-SIG ENGINEERING CHANGE REQUEST

TITLE:	~~Downstream Port Containment~~ _HPX and PCIe Completion Timeout related _OSC Enhancements
DATE:	Septem~~e~~ber 12, 2018
AFFECTED DOCUMENT:	PCI Firmware Specification, Rev. 3.2
SPONSOR:	Austin Bolen, Dell EMC

Part I

1. Summary of the Functional Changes

Changes are requested to be made to Section 4.5.1, _OSC Interface for PCI Host Bridge Devices and Section 4.5.2.4 Dependencies Between _OSC Control Bits.

The changes will enable the Operating System to advertise if it is capable of support _HPX PCI Express Descriptor Setting Record (Type 3) to the firmware. It also enables the Operating System and the Firmware to negotiate ownership of the PCIe Completion Timeout registers.

2. Benefits as a Result of the Changes

A new _HPX mechanism has been added to the ACPI Specification (_HPX PCI Express Descriptor Setting Record - Type 3). This new _HPX allows Firmware to provide platform-specific settings for PCI Express registers not covered by the older _HPX setting records. Firmware needs to know if the Operating System supports the new _HPX method. If not supported by the Operating System, then the Firmware will be able to take corrective action (programming the registers via alternate mechanisms, for instance).

In addition, there have been new registers added to the PCI Express Capability Structure since PCI Express Capability Structure control (bit 4 of _OSC) was defined. There are registers related to PCIe Completion Timeout that are typically programmed by Firmware but there is not mechanism for Firmware to grant control of the PCI Express Capability structure while maintaining control of PCIe Completion Timeout registers. This proposal will define such a mechanism.

3. Assessment of the Impact

Table 4-4, Table 4-5 and Table 4-6 in Section 4.5.1 must be extended to add _HPX and PCIe Completion Timeout Value related Support and Control bits. Section 4.5.2.4 must be updated to referenc the new PCIe Completion Timeout Value related _OSC bits.

4. Analysis of the Hardware Implications

No impact.

5. Analysis of the Software Implications

These software changes are optional.

The Operating System needs to set a new bit in _OSC if it support the new _HPX PCI Express Descriptor Setting Record (Type 3).

If the Operating System wants ownership of PCIe Completion Timeout Value registers, it needs to add support for requesting ownership of PCIe Completion Timeout Value registers independent of requesting access for PCI Express Capability Structure control. If the Operating System is currently modifying PCIe Completion Timeout Value registers, it will need to avoid do so if not granted control of these registers even if it has been granted control of the PCI Express Capability Structure.

6. Analysis of the C&I Test Implications

No impact.

Part II

Detailed Description of the change

Change Section 4.5.1 Table 4-4 as follows:

Table 4-4 Interpretation of _OSC Support Field

Support Field bit offset	Interpretation
..	..
8	**_HPX PCI Express Descriptor Setting Record (Type 3) Supported** The operating system sets this bit to 1 if it supports _HPX PCI Express Descriptor Setting Record (Type 3). See the ACPI specification for the definition of _HPX PCI Express Descriptor Setting Record (Type 3). Otherwise, the operating system sets this bit to 0.

Change Section 4.5.1 Table 4-5 as follows:

Table 4-5 Interpretation of _OSC Control Field, Passed in via Arg3

Support Field bit offset	Interpretation	Inter-dependencies
..	..	..
8	**PCI Express Completion Timeout Control** The operating system sets this bit to 1 to request control over the PCI Express Completion Timeout fields. These fields include Completion Timeout Value (Device Control 2 Register bits 3:0 in the PCI Express Capability Structure) and Completion Timeout Disable (Device Control 2 Register bit 4 in the PCI Express Capability Structure). If the operating system successfully receives control of this feature, it is responsible for configuring the PCI Express	Yes, refer to Section 4.5.2.4.

	Completion Timeout fields in a manner that complies with the PCI Express Base Specification. Additionally, the operating system is responsible for saving and restoring all PCI Express Completion Timeout field settings across power transitions to and from S1, S2, S3 system power states when register context may have been lost.	

Change Section 4.5.1 Table 4-6 as follows:

Table 4-6 Interpretation of _OSC Control Field, Returned Value

Support Field bit offset	Interpretation	Inter-dependencies
..	..	..
8	**PCI Express Completion Timeout Control** The firmware sets this bit to 1 to grant control over the PCI Express Completion Timeout fields. If the firmware does not grant control of this feature, firmware must handle configuration of the PCI Express Completion Timeout fields. If firmware grants the operating system control of this feature, any firmware configuration of the PCI Express Completion Timeout fields may be overwritten by an operating system configuration, depending on operating system policy.	Yes, refer to Section 4.5.2.4.

Change Section 4.5.2.4 as follows:

4.5.2.4. Dependencies Between _OSC Control Bits

Because handling of hot-plug events, power management events, ~~and~~ advanced error reporting, and completion timeout all require the modification of PCI Express Capability registers, the operating system is required to claim control over the PCI Express Capability (bit 4 of the Control field) in conjunction with claiming control over PCI Express Native Hot Plug, PCI Express Native Power Management Events, ~~or~~ PCI Express Advanced Error Reporting, or PCI Express Completion Timeout (bits 0, 2, ~~and~~ 3, and bit 8 of the Control field). If the operating system attempts to claim control of any of these features without also claiming control over the PCI Express Capability, the firmware is required to refuse control of the feature being illegally claimed and mask the corresponding bit.